



# **Life 3.0: Being Human in the Age of Artificial Intelligence**

*Max Tegmark*

[Download now](#)

[Read Online](#) ➔

# Life 3.0: Being Human in the Age of Artificial Intelligence

*Max Tegmark*

**Life 3.0: Being Human in the Age of Artificial Intelligence** Max Tegmark

*New York Times* Best Seller

**How will Artificial Intelligence affect crime, war, justice, jobs, society and our very sense of being human? The rise of AI has the potential to transform our future more than any other technology—and there's nobody better qualified or situated to explore that future than Max Tegmark, an MIT professor who's helped mainstream research on how to keep AI beneficial.**

How can we grow our prosperity through automation without leaving people lacking income or purpose? What career advice should we give today's kids? How can we make future AI systems more robust, so that they do what we want without crashing, malfunctioning or getting hacked? Should we fear an arms race in lethal autonomous weapons? Will machines eventually outsmart us at all tasks, replacing humans on the job market and perhaps altogether? Will AI help life flourish like never before or give us more power than we can handle?

What sort of future do *you* want? This book empowers you to join what may be the most important conversation of our time. It doesn't shy away from the full range of viewpoints or from the most controversial issues—from superintelligence to meaning, consciousness and the ultimate physical limits on life in the cosmos.

## **Life 3.0: Being Human in the Age of Artificial Intelligence Details**

### **Life 3.0: Being Human in the Age of Artificial Intelligence**

364 pages , Published August 29th 2017 by Knopf (first published August 23rd 2017)

 [Download Life 3.0: Being Human in the Age of Artificial Intellig ...pdf](#)

 [Read Online Life 3.0: Being Human in the Age of Artificial Intell ...pdf](#)

**Download and Read Free Online Life 3.0: Being Human in the Age of Artificial Intelligence** Max Tegmark

---

## Manny says

Last week, I read Nick Bostrom's groundbreaking *Superintelligence*, an extremely serious, well thought out look at the dangers of creating real artificial intelligences. It left me feeling more than a little concerned: despite working in AI myself, I had not fully appreciated how scary it is. I've just finished this book, written about three years after Bostrom's, and now I'm even more concerned.

Bostrom's book has the air of being primarily intended for senior policy-makers in industry and government. It's dry and formal in tone, relentlessly footnoted, full of difficult words, and seldom goes out its way to try and entertain you. It appears to be the product of a great deal of diligent work, and the dominant note is one of a dire, Cassandra-like warning. But despite the fact that it quotes Bostrom on every other page, Tegmark's book is pretty much the opposite: it's manic and chatty, dumbed down to one of the lower common denominators, poorly structured, and full of winsome autobiographical revelations and engaging little science-fiction stories. I could have taken all that - I've read Tegmark's previous book, and I'm familiar with his style - but what really worries me is that he's trying to make us feel *optimistic* and *hopeful*. It'll be okay! He's founded his Future of Life Institute, all sorts of prominent AI people have joined up, they're organising conferences on AI safety and an incredible ongoing web conversation which anybody can join, and pretty soon they'll have the answer!

Well, if you believe that you'll believe anything. As Bostrom explains, the problems here are not technical issues that can be addressed by good engineering. They are foundational philosophical questions that have been discussed for thousands of years without a great deal of progress having been made. It's readily conceivable that there *are* no solutions to them, and there's nothing available right now which even looks vaguely like a possible road towards a solution. Anyone expecting a magic fix to turn up over the next couple of decades is probably also expecting it to be delivered in person by the Easter Bunny riding a pink fluffy unicorn. And yet Tegmark goes on and on about his preferred future, where we become spacefaring AIs who head out, à la Olaf Stapledon, to transform the universe into one huge superintelligence which, after several billion years of development, will be able to perform ten to some very big number of computations per second. He spends inordinately many pages explaining various tweaks, all of which will no doubt be invalidated by the next major discovery in physics, that mean that it will be ten to this big number rather than ten to some slightly smaller big number. Why we're supposed to do it is never once explained. Stapledon's Cosmic Mind gets to make contact with the Star Maker, which at least is a worthy goal, but what Tegmark is hoping for is beyond me. I guess it's just, you know, cool. But all the same, we're supposed to stake the whole future of humanity on this bizarre and geeky dream.

If Tegmark were a moron, I would be slightly reassured. I figured out some time ago that many books are written by morons. But he's a very clever person, who, at least on his own account, has considerable influence over the Secret International Network of Very Clever People. I suppose he may be fronting the superintelligent AI who's currently taking over the world; I don't know if that would make it better or worse, but at least it would make it more comprehensible. Aaaargh!

To conclude, I can't do better than to quote a couple of passages which frequently occurred to me while reading *Life 3.0*. The first, taken from the final section of Bostrom's book, is one conspicuously not cited by Tegmark:

Before the prospect of an intelligence explosion, we humans are like small children playing with a bomb. Such is the mismatch between the power of our plaything and the immaturity of our conduct. Superintelligence is a challenge for which we are not ready now and will not be ready for a long time. We have little idea when the detonation will occur, though if we hold the device to our ear we can hear a faint ticking sound.

For a child with an undetonated bomb in its hands, a sensible thing to do would be to put it down gently, quickly back out of the room, and contact the nearest adult. Yet what we have here is not one child but many, each with access to an independent trigger mechanism. The chances that we will *all* find the sense to put down the dangerous stuff seem almost negligible. Some little idiot is bound to press the ignite button just to see what happens.

The second passage, which has been one of my favorites since I was a teen, is from another exceptionally clear-sighted book, Kurt Vonnegut's *Cat's Cradle*:

And I remembered *The Fourteenth Book of Bokonon*, which I had read in its entirety the night before. *The Fourteenth Book* is entitled 'What Can a Thoughtful Man Hope for Mankind on Earth, Given the Experiences of the Past Million Years?'

It doesn't take long to read *The Fourteenth Book*. It consists of one word and a period.

This is it:

'Nothing.'

---

[And on further consideration...]

As you can see, I found *Life 3.0* extremely annoying. But at the same time, I also appreciate the point Tegmark makes, that it doesn't help just to say that we're all going to die. Some more positive approach is required.

It seems to me, given the analysis in Bostrom's very sensible book, that there is little chance of making superintelligence 'safe' in any normal sense of the word. If we create a superintelligence, it will almost certainly replace us before long. So it follows that we have two realistic choices:

- a) Do not create a superintelligence.
- b) Create a superintelligence that we could happily envisage replacing us.

I think most people will go for (a). If we are even going to consider (b), I would say that the priority should be to develop AIs with qualities that people will value, and which they feel are characteristic of what it truly means to be human. Some obvious candidates include creativity, humour, curiosity, honesty, appreciation of beauty, empathy and love.

Tegmark briefly mentions "creativity", though his analysis of the already renowned move 37 in the second AlphaGo-Lee Sedol game is as superficial as the rest of the book. I see no real discussion of the other qualities.

### **Brian Clegg says**

I have to confess that my first reaction to this book was not anything to do with the contents, but trying to work out if there was something really clever about the way the book's title is printed on the spine in white on cream, so it's illegible - would it be, for example, a subtle test of human versus artificial intelligence (AI)? However, that was just a distraction.

Max Tegmark is an interesting and provocative thinker in the physics arena, so I had high hopes for what he'd come up with exploring the future of AI and its relationship to human beings. It's worth explaining that the title of the book refers to three 'levels' of life where 1.0 is 'can survive and replicate' (e.g. bacteria), 2.0 is can design its own software (e.g. us - where 'software' refers to our concepts, ideas and extended abilities such as language) and 3.0 is can design its own hardware, enabling it to transform itself more directly and quickly than our creativity enables us to do.

The book starts with a bit of fiction, which I'm usually nervous about, but it actually works very well, as it's presented more like a non-fiction description of a business development rather than attempting all the quirks of fiction. In it we have a semi-plausible description of how a company that succeeds in producing a self-enhancing AI could take over the world. And this is genuinely thought-provoking.

So, early on, I was convinced I was going to love this book. But unfortunately there is an awful lot of futurology in here (aka guesswork) and like all futurology, Tegmark's can be frustratingly specific about things that we are highly unlikely to be able to predict - though at least he recognises this is the case and points it out. He covers the various ways a super-intelligent AI could develop, whether it would become a rogue, how we'd interact with it... and then plunges on into more and more dramatic speculation, including a chapter that looks forward 'the next billion years and more.' Forgive me for feeling a bit 'So what?' about this.

There is no doubt the whole business of super-AIs is an issue that needs thinking about and discussing - and Tegmark does this in an approachable and engaging fashion. It probably needs reading alongside Nick Bostrom's *Superintelligence* to get a well-rounded picture, though. It would have helped if it had been significantly shorter - it came across as being long because it was the kind of 'big book' that has to be chunky, and I think it would have been a lot more effective at half its length. One particular section that was ripe for trimming had a long list of scenarios, each of which was then worked through - dull reading, I'm afraid.

In the end futurology is a bit like being told about someone else's dream. It probably seems fascinating to them, but it's hard to get too excited about it as a reader. *Life 3.0* is an interesting book, but feels rather like a pet project, rather than a strong popular science title.

### **Lor-El says**

Unlike the author, I will try to be concise and make my point clear:

1) There is so much name-dropping (authors, books, theories, Ivy League Universities, tv shows, movies... you name it, it's gonna be there) and basically it seems like a secondary literature review rather than an original work.

2) Some chapters feels like fillers, put there just to make the book thicker, they add little to no useful information on Humans and AIs whatsoever.

3) The long awaited (like, 300+ pages awaited) chapters on all the ethical questions and consciousness, which would have made the book worth reading for me, are just a scratch on the surface and do not delve into the depth of these multilayered issues. A show like "Person of Interest" (which, although the Prometheus-AI takeover described in the book is uncannily similar to the one perpetrated by Samaritan, is not even once cited in the book) did a better job at dealing with both ethics, machine learning and how human deal with two sentient and very different ASI. And I know... I know, "that's just sci-fi" but ultimately, so is this book.

# **From Reader Review Life 3.0: Being Human in the Age of Artificial Intelligence for online ebook**

**From reader reviews:**

**William Reeves:**

Why don't make it to be your habit? Right now, try to prepare your time to do the important act, like looking for your favorite reserve and reading a book. Beside you can solve your problem; you can add your knowledge by the reserve entitled Life 3.0: Being Human in the Age of Artificial Intelligence. Try to make the book Life 3.0: Being Human in the Age of Artificial Intelligence as your friend. It means that it can to be your friend when you experience alone and beside that of course make you smarter than ever. Yeah, it is very fortunated in your case. The book makes you much more confidence because you can know every thing by the book. So , we should make new experience and also knowledge with this book.

**Lisa Maurer:**

This Life 3.0: Being Human in the Age of Artificial Intelligence book is not ordinary book, you have it then the world is in your hands. The benefit you receive by reading this book is definitely information inside this reserve incredible fresh, you will get facts which is getting deeper you actually read a lot of information you will get. This particular Life 3.0: Being Human in the Age of Artificial Intelligence without we recognize teach the one who examining it become critical in contemplating and analyzing. Don't be worry Life 3.0: Being Human in the Age of Artificial Intelligence can bring if you are and not make your carrier space or bookshelves' become full because you can have it with your lovely laptop even mobile phone. This Life 3.0: Being Human in the Age of Artificial Intelligence having fine arrangement in word along with layout, so you will not experience uninterested in reading.

**Marcos Gorman:**

The reason why? Because this Life 3.0: Being Human in the Age of Artificial Intelligence is an unordinary book that the inside of the book waiting for you to snap this but latter it will distress you with the secret it inside. Reading this book beside it was fantastic author who also write the book in such awesome way makes the content inside of easier to understand, entertaining technique but still convey the meaning fully. So , it is good for you for not hesitating having this any more or you going to regret it. This amazing book will give you a lot of advantages than the other book include such as help improving your proficiency and your critical thinking means. So , still want to hold up having that book? If I ended up you I will go to the publication store hurriedly.

**Mary Barrientes:**

The book untitled Life 3.0: Being Human in the Age of Artificial Intelligence contain a lot of information on the idea. The writer explains the woman idea with easy technique. The language is very clear to see all the people, so do not necessarily worry, you can easy to read this. The book was compiled by famous author. The author provides you in the new age of literary works. You can easily read this book because you can keep reading your smart phone, or device, so you can read the book within anywhere and anytime. In a situation you wish to purchase the e-book, you can wide open their official web-site in addition to order it.

Have a nice examine.

**[BOOK]? Life 3.0: Being Human in the Age of Artificial Intelligence Max Tegmark WYX3VT6ZEBN**